

Achieving Fair Network Equilibria with Delay-based Congestion Control Algorithms

Miguel Rodríguez-Pérez, Sergio Herrería-Alonso, Manuel Fernández-Veiga, *Member, IEEE*,
Andrés Suárez-González and Cándido López-García

Abstract—Delay-based congestion control algorithms provide higher throughput and stability than traditional loss-based AIMD algorithms, but they are inherently unfair against older connections when the queuing and the propagation delay cannot be measured accurately and independently. This paper presents a novel measurement algorithm whereby fairness between old and new connections is preserved. The algorithm does not modify the dynamics of congestion control, and runs entirely in the server host using locally available information.

Index Terms—Delay-based congestion control, FAST TCP, persistent congestion, fairness.

I. INTRODUCTION

DELAY-BASED congestion avoidance (DCA) algorithms, such as FAST or Vegas, achieve high throughput in high-speed long-latency networks [1], [2]. But it is also well known that their equilibrium transmission rates are very sensitive both to the accuracy of the estimated round-trip propagation delay and to the estimated queuing delay. Measurement errors in any of these quantities may lead to severe unfairness. A situation like that arises, for instance, when a new flow encounters a state where the queue ahead of the bottleneck link never gets empty, thus hampering to correctly estimate the propagation delay along its network path. This harmful, self-sustained condition, termed *persistent congestion*, was already found as early as in [3].

In [4], a mathematical analysis is provided for a scenario where persistent congestion is due to the successive arrival of a set of everlasting flows to an empty router queue. It has been argued that such scenario is far unlikely, however, the arrival of just a single flow to a saturated link is a sufficient condition to trigger unfairness as long as some of the older flows do not depart. Such configuration, where a small group of newborn flows find a link in equilibrium (bandwidth equally distributed) shared by n preexisting long-lived flows, was precisely the setting analyzed in [5], and includes [4], in fact, as a particular case. As a possible solution to the persistent congestion problem, [5] suggests throttling down briefly each newly started flow to allow queues to empty, and thus obtain a reliable estimate of the propagation delay. We have found that this approach is not always effective, though.

We show that such a cautious source can fail to measure a correct propagation delay under general circumstances, and

present a novel solution able to remove the undesired effect of persistent congestion in arbitrary conditions. As in [5], our proposal only requires the modification of the sender end host, and attains a throughput as high as (and a buffer utilization as low as) FAST does.

II. EQUILIBRIUM RATE OF RECENT ARRIVALS

Despite their differences at the packet level, all congestion control algorithms can be mathematically described, at the flow level, by the dynamical equation

$$\dot{w}_i(t) = \kappa_i(t) \left(1 - \frac{q_i(t)}{u_i(t)} \right) \quad (1)$$

where $w_i(t)$ denotes the congestion window at time t for flow i , $\kappa_i(t)$ is a gain function, $u_i(t)$ is a suitable utility function, and $q_i(t)$ is the congestion signal [1]. The transmission rate is then given by $x_i(t) = w_i(t)/r_i(t)$, where $r_i(t)$ is the round-trip time. For DCA algorithms, $q_i(t)$ is the queuing delay. TCP Vegas uses $\kappa_i(t) = 1/r_i(t)$, whereas FAST takes $\kappa_i(t) = \gamma\alpha/\tau$, where γ , α and τ are protocol parameters. Both instances, FAST and Vegas, use $u_i(t) = \alpha/x_i(t)$ and have therefore equal equilibrium structure, determined by (1), namely

$$x_i^* = \frac{\alpha}{r_i^* - \hat{d}_i}, \quad (2)$$

where \hat{d}_i is the propagation delay as estimated by flow i .

We consider in this paper the arrival of a *single* new flow (indexed by 0) at a bottleneck link of capacity C shared by a set $\mathcal{F} = \{1, \dots, n\}$ of FAST flows. We also assume that each connection $f \in \mathcal{F}$ knows its true round-trip propagation delay ($\hat{d}_f = d_f$).¹ Hence, each flow f is receiving C/n units of bandwidth. Following the model in [5], flow f contributes α packets to the router queues, so flow 0 sees a propagation delay of $\hat{d}_0 = d_0 + n\alpha/C$. As a result of this overestimated value, it grabs a rate in the equilibrium

$$x_0^* = \frac{\alpha}{r_0^* - \hat{d}_0} = \frac{\alpha C}{b_0^*}, \quad (3)$$

while the new common equilibrium rate for the older flows is

$$x_f^* = \frac{\alpha C}{b_0^* + n\alpha}. \quad (4)$$

Since $\sum_{i=0}^n x_i^* = C$ we obtain

$$b_0^* = \frac{\alpha}{2} (1 + \sqrt{1 + 4n}). \quad (5)$$

The authors are with the Telematics Engineering Dept., Univ. of Vigo, 36310 Vigo, Spain. Tel.: +34 986 813459; fax: +34 986 812116; email: miguel@det.uvigo.es (M. Rodríguez-Pérez). This work was supported by the “Ministerio de Educación y Ciencia” through the project TS12006-12507-C03-02 of the “Plan Nacional de I+D+I” (partly financed with FEDER funds).

¹That is, we assume that there is a working algorithm in place to account for the persistent congestion bias. In Section IV we present such an algorithm.

The transmission rates given by (3) and (4) are clearly unfair in that the recent arrival obtains far more bandwidth than the rest. Moreover, the unfairness worsens with the number of flows, $x_0^*/x_f^* \sim O(\sqrt{n})$.

We claim that the fair equilibrium is achievable using a slightly modified procedure to measure the propagation delay (see Section IV). Hence, since the onset of persistent congestion can be completely avoided, any new flow will find the bottleneck link capacity fully and equally shared among the older ones, as long as their rates have stabilized during the time elapsed from the last arrival. Consequently, there is no need to pose the case of successive flow beginnings, as in [5], and the assumption of a single recent arrival does not entail loss of generality.

III. THE RATE REDUCTION APPROACH

The solution presented in [5] consists in restraining transiently the transmission rate of a new flow by a given factor to allow router queues to get eventually empty, thus giving new connections a chance to directly measure the true round-trip propagation delay. Unfortunately, and despite of the reduction on its rate, the new connection is not always able to detect queue emptiness. Note that, as the new flow drains queues by reducing its own rate, competing flows respond by increasing their rates. Hence, the new flow will only obtain the true propagation delay if queues empty before existing flows are aware of this event, that is, if the time required to empty the queues is less than the RTT of the existing flows.

Let $B^* = b_0^* + n\alpha$ be the total backlog buffered at the core of the network in equilibrium. This backlog will be drained from the queue at a rate equal to the bottleneck link capacity minus the sum of the transmission rates of all active flows. In the most favorable case, the new connection will completely pause its transmission ($x_0 = 0$). Then, if all the existing flows $f \in \mathcal{F}$ experience the same propagation delay ($d_f = d$), and so the same RTT ($r_f^* = r^*$), the fairness condition becomes

$$\frac{B^*}{C - \sum_{f=1}^n x_f^*} < r^* = d + \frac{b_0^* + n\alpha}{C}. \quad (6)$$

Finally, substituting (5), (4) and B^* into (6), it follows that

$$d > \frac{n\alpha(1 + \sqrt{1 + 4n})}{2C} = \frac{nb_0^*}{C} \quad (7)$$

Thus, the rate reduction method is only effective when the round-trip propagation delay of competing flows exceeds the lower bound calculated in (7). This lower bound scales as $O(n^{3/2})$ with the number of active flows, preventing a sensible default for the duration of the rate reduction.

IV. A NOVEL SOLUTION

We noticed that, when the newly arriving flow stabilizes, it can indirectly obtain a good estimation of its actual round-trip propagation delay. As already pointed in Section II, the new flow overestimates its propagation delay as $\hat{d}_0 = d_0 + n\alpha/C$. Since \hat{d}_0 and α are known, it suffices to estimate n/C to get the real d_0 .

A good estimation of n and C can be obtained even if the router queues are not completely empty. In fact, as we

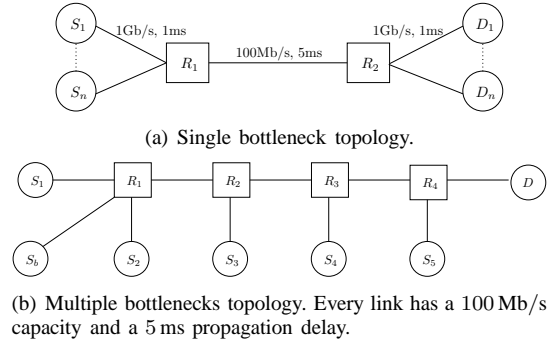


Figure 1. Network topologies used in the simulation experiments.

will show, it suffices to just indirectly measure queue length variations after a short change of the transmission rate. Let r_0^* be the RTT of the tagged flow once it reaches a stable throughput. If this connection modifies its transmission rate $x'_0 = (1 - \theta)x_0^*$, with $\theta < 1$, for a brief time t_ϵ (of the same order as r^* , so that the rest of the flows do not adjust their own transmission rates) it will measure a new RTT r'_0 when it resumes its transmission. Let $\Delta r_0 = r_0^* - r'_0$. Under such circumstances

$$C\Delta r_0 = \left(C - \sum_{i=1}^n x_i^* - (1 - \theta)x_0^* \right) t_\epsilon \quad (8)$$

Substituting (3), (4) and (5) in (8), and solving for n yields

$$\hat{n} = \frac{\theta t_\epsilon}{\Delta r_0} \left(\frac{\theta t_\epsilon}{\Delta r_0} - 1 \right). \quad (9)$$

Now, using (9) and (3) $\hat{C} = \frac{(1 + \sqrt{1 + 4\hat{n}})x_0^*}{2}$ and the correct propagation delay can be adjusted as

$$\hat{d}'_0 = \hat{d}_0 - \alpha \frac{\hat{n}}{\hat{C}}.$$

Note that using positives values for θ causes the queue to drain, and it is possible to exhaust the backlog before the end of the measure. In that case (8) no longer holds and the number of flows is overestimated. To avoid it, it suffices to use small negative values for θ , causing the queueing delay to increase. Although for insufficiently dimensioned buffers this may cause some packet drops, this condition can be easily detected and avoided by using smaller values of θ in subsequent measures.

V. PERFORMANCE ANALYSIS

To verify these claims, we report several ns-2 simulation experiments. In the first one, there are five FAST connections (S_i, D_i) sharing the bottleneck link (Fig. 1(a)), starting at intervals of 20 s each. Routers' buffers are large enough to avoid packet losses, and sources always have data to send.

Fig. 2(a) shows the instantaneous throughputs of the FAST flows with the original congestion avoidance mechanism ($\alpha = 50$ packets). As expected, FAST strongly favors new sources and recent connections get larger throughput than older flows. With the modified measurement method, this bias disappears and the network bandwidth is shared fairly (Fig. 2(b)).² Also,

²For the estimation of \hat{n} we have employed $t_\epsilon = r^*$ and $\theta = -0.5$ to prevent the bottleneck from getting empty.

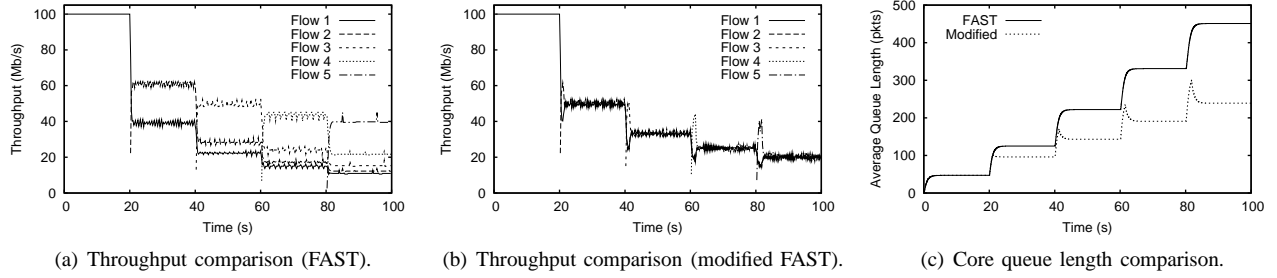


Figure 2. Throughput and core queue length comparisons.

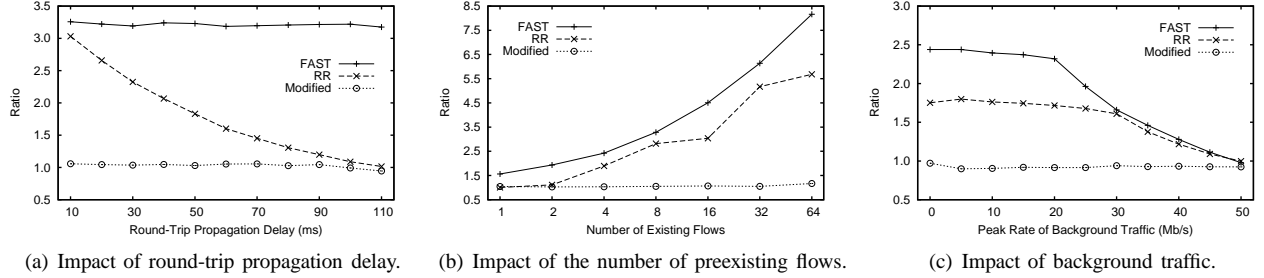


Figure 3. Simulation experiment results.

the average queue length at the bottleneck (Fig. 2(c)) is consistently lower because, due to persistent congestion, the backlog of FAST exceeds the target value of α packets per source, whereas our proposal does not so.

A second test was run over the same network to compare the proposed algorithm with the original FAST protocol and the rate reduction (RR) variant. Assume a set of existing FAST flows aware of their true propagation delays, sharing the bandwidth uniformly. Once their rates stabilize, a new flow starts. The delay of the link (R_1, R_2) was appropriately set so as to have the desired RTT. Following customary practice, we measured the fairness among the new and the n existing connections as the ratio $n\bar{x}_0 / \sum_{f=1}^n \bar{x}_f$ where \bar{x}_0 is the average transmission rate of the new flow and \bar{x}_f denotes the average rate of flow $f = 1, \dots, n$. Fig. 3(a) compares the performances of the three protocols for $n = 8$. As expected, with FAST, the new connection obtains a higher throughput. With the RR method, the bandwidth sharing depends on the experienced propagation delay: for delays below the threshold given by (7) (108 ms in this scenario), the source rates become unfair. In contrast, fairness is preserved if the novel solution is used. Further, for any given RTT, the unfairness aggravates with the number of flows, as Fig. 3(b) clearly shows, either for FAST or for the RR reduction method, and only the modified version allocates bandwidth equally.

A more realistic and stringent topology was also considered. In Fig. 1(b), the network (a variant of the classic parking-lot topology) has multiple bottlenecks, with five flows running from nodes S_1, \dots, S_5 to node D . The flow originated in S_1 starts its transmission after the rest of the flows stabilize. Additionally, in a similar way as in [1], some background traffic was simulated with a Pareto flow (S_b, D) with shape factor of 1.25, average burst and idle time of 100 ms and a peak rate ranging from 5 to 50 Mb/s. Fig. 3(c) shows the results.

Not surprisingly, with both FAST and the RR method, fairness improves as the peak rate of background traffic increases. The reason is that, during active periods, FAST flows reduce their rates as router queues fill due to background traffic, so in the idle periods the new flow can seize a better estimate of its propagation delay before queues fill again. In any case, our solution assures fairness irrespective of the amount of background traffic introduced.

VI. CONCLUSIONS

This paper has demonstrated that the *rate reduction approach* fails to solve persistent congestion in networks shared by many flows, as it cannot always completely drain the bottleneck queues, and thus is unable to obtain an accurate measure of the propagation delay.

We have presented a novel solution that does not rely on getting a direct measure of the propagation delay. Instead, by carefully modulating its own transmission rate, the source is able to calculate the error in the estimation of the round trip propagation delay and thus share the link evenly with the other FAST flows onwards.

REFERENCES

- [1] D. X. Wei, C. Jin, S. H. Low, and S. Hedge, "FAST TCP: Motivation, architecture, algorithms, performance," *IEEE/ACM Trans. Networking*, vol. 14, no. 6, pp. 1246–1259, Dec. 2006.
- [2] L. Brakmo, S. O'Malley, and L. Peterson, "TCP Vegas: New techniques for congestion detection and avoidance," in *Pr. SIGCOMM'94*, 1994, pp. 24–35.
- [3] J. Mo, R. La, V. Anantharam, and J. Walrand, "Analysis and comparison of TCP Reno and Vegas," in *Proc. INFOCOM'99*, 1999, pp. 1556–1563.
- [4] S. H. Low, L. Peterson, and L. Wang, "Understanding Vegas: a duality model," *J. ACM*, vol. 49, no. 2, pp. 207–235, Mar. 2002.
- [5] T. Cui, L. Andrew, M. Zukerman, and L. Tan, "Improving the fairness of FAST TCP to new flows," *IEEE Commun. Lett.*, vol. 10, no. 5, pp. 414–416, May 2006.